

Understanding the Role of Automated Response Actions in Improving AMI Resiliency

Ahmed Fawaz, Robin Berthier, William H. Sanders
Information Trust Institute and
Department of Electrical & Computer Eng.
University of Illinois at Urbana-Champaign
Urbana, IL, USA
Email: {afawaz2, rgb, whs} @illinois.edu

Partha Pal
BBN Technologies
Cambridge, MA, USA
Email: ppal@bbn.com

Abstract—The smart grid promises better services and higher reliability but is exposed to new security threats. In particular, deployment of advanced metering infrastructures (AMIs) will vastly increase the attack surface because of the smart meters installed in customer homes. Managing the security of AMIs cannot be done manually because of their size and complexity. Thus, we propose a three-step plan to bring automated responses to AMIs. Considering the challenges of automated responses, we will develop a taxonomy of response actions in AMIs. Then, we will model the response actions in terms of their impact and cost for the different actors in the system: users, administrators, and attackers. Finally, we will discuss implementation and evaluation requirements for a practical automated response engine for AMIs.

Index Terms—AMI, CPS, Response action, Cyber security.

The adoption and deployment of the smart grid promise customers faster and more reliable service. The smart grid is enabling those improvements through its capabilities for remote control, instant detection of blackouts, and accurate state estimation of the power grid using phasor measurement units (PMU). Additionally, the smart grid accommodates more customer services, such as real-time pricing, and includes provisions for future electrical vehicles. Advanced metering infrastructures (AMI) are a core component of the smart grid that is now being deployed. AMIs are the communication solution for smart meters that transmit real-time meter readings to the administrative network and receive commands to control service remotely. AMIs enable fine-grained detection of blackouts and will thus enable faster customer service.

A typical AMI will allow remote control of every smart meter, including the ability to turn service off, and in some scenarios will allow utility companies to control specific appliances in individual homes as part of environmental programs that offer reduced prices at certain hours of the day. Moreover, utility companies will no longer need to have human meter readers drive around and obtain monthly readings, because readings will be sent to the utility company frequently from the meters through the AMI network. Finally, the introduction of smart appliances that can communicate with smart meters to get real-time pricing information means that owners will be able to control those appliances remotely via the Internet.

AMIs present more security problems than regular cyber-physical systems (CPS) do, as their architecture and services

allow for a larger attack surface. The attack surface includes 1) the corporate network, 2) the wireless mesh network, 3) the home area network, and 4) meters that are within the reach of customers. Possible threats can be classified according to attack scale, ranging from relatively small-scale targeting of specific houses (in order to turn off service or specific appliances, such as alarm systems) or stealing of energy (through alteration of meter readings or duplication of meters), up to large organized crimes that target large geographical regions. Moreover, attacks could target the control commands sent by a utility to its AMI. Additional security issues also rise from the use of the wireless channel for smart meter communication. Additional attacks will be facilitated by the wireless mesh network that will be used to connect meters; such networks are prone to single points of failure, availability problems, jamming, eavesdropping, man-in-the-middle attacks, and wormhole and black hole attacks [1], [3], [5].

Compared to traditional IT systems, AMIs have stringent requirements in terms of quality of service and security guarantees. Those requirements include:

- 1) *Availability*: Utility companies should be able to get the latest meter readings and send out control commands within specific time constraints. Moreover, customers expect the latest pricing to be available.
- 2) *Resilience*: AMI provides a critical service to customers. It must be able to work under extreme conditions and provide the core service of measuring energy consumption even under attack.
- 3) *Fast recovery*: In the event of an attack, a compromise, equipment faults, or even blackouts, an AMI should allow fast recovery and restoration of service.
- 4) *Size*: In the future, a typical AMI could be larger than any conventional CPS ever built, with millions of nodes in cities; this massive size imposes scalability issues for traditional security solutions.
- 5) *Privacy*: There are also privacy concerns specific to AMIs, since the readings and commands sent between the meter and the utility company reveal private information about customers.

Important efforts (by researchers and by organizations such

as NERC and NIST) have been made to promote security solutions for AMI networks, such as VPNs, encryption [4], and remote attestation [8]. Such efforts are important, but cannot completely secure systems, mainly because vulnerabilities can always be found in the implementations of protocols and applications, or in the human operators who can be tricked into providing access to restricted resources. Moreover, since meters are left without real physical protection, tampering with devices may leak secret keys stored in internal memory and thus cause security breaches in the network. Thus, traditional attack prevention solutions have to be supplemented with detection and mitigation approaches. Our present work focuses on studying the possibility of a framework that can automatically respond to cyber intrusions given the requirements of an AMI.

The importance of intrusion detection for AMIs is still critical, and several approaches have been proposed [2], [11]. However, intrusion detection is prone to inaccuracies, and monitoring such a large number of nodes will rapidly lead to an unmanageable volume of alerts and demands for decisions. The combination of potentially weak detection capabilities and stringent CPS requirements means that to offer strong resiliency against cyber-attacks, security solutions have to be proactive. For example, the uncertain identification of a suspicious behavior has to trigger the automated deployment of additional monitoring capabilities to translate inaccurate reports into actionable information. A variety of automated response solutions have been studied over the past decade [13], but none have been tailored for the specific requirements of complex cyber-physical systems such as AMIs. Moreover, the practicality of existing solutions is limited, and for multiple reasons, the industry has been reluctant to implement sophisticated automated response actions. First, implemented actions are often all-or-nothing, meaning that they lack the flexibility to adapt to various situations and can lead to dramatic consequences in the case of false positives. Second, we have a poor understanding of the impact of response actions in large and complex systems. Third, that lack of understanding can result in vulnerabilities in the response action itself, which could enable attackers to game the system and cause automation to do more harm than good.

We gained a better understanding of the limitations of current automated response solutions by reviewing related work from the perspective of practicality for the specific requirements of AMIs. As a result, we plan to present the following approach to bringing efficient and secure automated response to AMIs.

The first step, which is in progress, involves development of a taxonomy of response actions that suits AMI requirements, such as always preserving the mission of delivering energy and accurately measuring consumption. The taxonomy will allow us to construct a set of possible response actions by emphasizing the concept of flexibility. Flexible actions can be tuned to meet a wide variety of requirements and situations. This will then guide the development of a practical case study of ways to design flexible actions for an AMI. The taxonomy

has two high-level categories: 1) *learning* actions, and 2) *modifying* actions. Learning actions are either passive or active and are designed to gather additional information about security incidents. Learning actions include enabling of additional IDS sensors with a higher granularity, logging of traffic, or active sending of probe packets to locate compromised nodes. Modifying actions work to respond to and recover from an attack. Modifying actions have two subcategories: limiting actions and recovery actions. Limiting actions reduce privileges of a given entity, thus reducing its ability to propagate an attack. Limiting actions include addition of firewall rules to block a meter's traffic, changes to access privileges to certain resources within a meter, and changes to routes within the mesh network to avoid a compromised meter. Recovery actions will work to stop attacks and return to a previous working state in the system; such responses include application of update patches, flashing of a clean OS version, and even sending of field technicians to change a meter.

The second step after building the taxonomy will be to model the response actions' impact and cost. The first task in modeling response actions will be to identify the different actors in our system. Usually, security researchers consider the main actors to be the *attacker* and the *administrator*. However, we propose to include customers as well, since they can also be affected by the attacker's actions and the administrator's reactions. An action's impact can be described as *beneficial* or *harmful*, where *beneficial* actions are those that benefit legitimate entities (administrators and customers) and negatively impact illegitimate entities (i.e., attackers) by making it harder for them to achieve their malicious goals. In order to quantify the impact of an action, it is necessary to define the cost of the action. Several researchers have proposed methods to compute the cost of actions [6], [7], [10]; some propagate an availability metric after an action is. Others decompose the cost based on the number of unavailable resources, impact on the system, and operation cost. Most research uses a weight matrix for the different confidentiality, integrity, and availability (CIA) metrics to describe the importance of each security property. Most previous work does not look into practical ways to compute the cost of an action to the user, or consider the time needed to recover as part of the cost. Moreover, the use of a static matrix to specify the importance of each security property is highly subjective and does not provide a method to compute those values based on the policies of the corporation, or even provide a sense of how to tweak the values to change the reactions of the system. Additionally, cost assessment in the context of a CPS requires a detailed understanding of the interfaces between the cyber and physical mechanisms. Because of those limitations, we will propose a cost computation method that allows us to consider the cost for customers. It will also allow for flexible cost for actions with varying intensity (e.g., rate limiting with a variable threshold rate). Moreover, we will propose clear methods to generate and tweak the weight coefficients needed to compute the cost of an action, as well as include the impact on physical systems in the calculation.

The next step in this project will be to explore solutions for automatic selection of response actions at runtime during an attack. We plan to study the game-theoretic response and recovery engine (RRE) proposed by Zonouz et al. [12]. RRE models the system as a Stackelberg game [9] between the attacker and the administrator. RRE uses an attack-response tree to represent the possible attacker moves and tags each move with a set of possible responses. Upon an attacker's move, RRE then computes an optimal strategy for the current security state of the system that maximizes the benefit for the administrator while reducing the benefit for the attacker. Several challenges must be addressed before such online automated decision-makers become "AMI-ready." First of all, because of the large size of an AMI, the attack-response tree representing the system will get much larger than those of traditional networks, making it difficult for RRE to compute the optimization. Thus, we need an abstraction to reduce the search space of RRE for the AMI. The main idea behind the abstraction is to use the hierarchy within AMIs, so we will divide the attack goals into several interim goals that can be solved independently within a neighborhood. Then, we will form another tree that combines several neighborhoods and decides on high-level actions (e.g., isolating a complete neighborhood). Moreover, RRE does not have provisions for customer costs, and changes are needed to include those costs as part of computing the optimal response strategy.

The final contribution of this project will be to discuss how to evaluate the framework in a realistic environment. We will present a set of experiments that we plan to implement in the TCIPG/Itron testbed. This testbed emulates hundreds of virtualized meters combined with hardware meters, all clustered into several neighborhoods (reflecting a realistic AMI). Each cluster has a collector that sends the readings back to the head end or sends commands from the head end to the meters.

This paper presents a rigorous research plan to study automated response within the unique requirements of an AMI. The proposed solution will help utilities improve on services, operation costs, and reliability. Automated response in AMIs will reduce the maintenance cost for utilities, as it will improve the ability to troubleshoot the distribution network by providing situational awareness. Moreover, automated

response has the potential to significantly reduce the load on human operators by automatically managing low-level alarms generated by sensors in the network.

ACKNOWLEDGEMENT

This material is based upon work supported by the Department of Energy under Award Number DE-OE0000097.

REFERENCES

- [1] M. Al-Shurman, S. M. Yoo, and S. Park, "Black hole attack in mobile ad hoc networks," in *Proceedings of the 42nd Annual Southeast Regional Conference*, 2004, pp. 96-97.
- [2] R. Berthier, W. H. Sanders, and H. Khurana, "Intrusion detection for advanced metering infrastructures: Requirements and architectural directions," in *Proc. First IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2010, pp. 350-355.
- [3] L. Buttyan and J.-P. Hubaux, *Security and Cooperation in Wireless Networks*. Cambridge University Press, 2007.
- [4] A. Hahn and M. Govindarasu, "Cyber attack exposure evaluation framework for the smart grid," *IEEE Transactions on Smart Grid*, vol. 2(4), pp. 835-843, 2011.
- [5] Y. C. Hu, A. Perrig, and D. B. Johnson, "Wormhole attacks in wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, pp. 370-380, 2006.
- [6] W. Kanoun, N. Cuppens-Bouahia, F. Cuppens, S. Dubus, and A. Martin, "Success likelihood of ongoing attacks for intrusion detection and response systems," in *Proc. International Conference on Computational Science and Engineering (CSE'09)*, 2009, pp. 83-91.
- [7] N. Kheir, H. Debar, N. Cuppens-Bouahia, F. Cuppens, and J. Viinikka, "Cost evaluation for intrusion response using dependency graphs," in *Proc. International Conference on Network and Service Security (N2S'09)*, 2009, pp. 1-6.
- [8] M. LeMay and C. Gunter, "Cumulative attestation kernels for embedded systems," in *Proceedings of the 14th European Conference on Research in Computer Security (ESORICS 2009)*, pp. 655-670, 2009.
- [9] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus, "Playing games for security: An efficient exact algorithm for solving Bayesian Stackelberg games," in *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*, vol. 2, 2008, pp. 895-902.
- [10] C. Strasburg, N. Stakhanova, S. Basu, and J. Wong, *The Methodology for Evaluating Response Cost for Intrusion Response Systems*, Technical Report 08-12, Iowa State University, 2008.
- [11] Y. Zhang, L. Wang, W. Sun, R. C. Green, and M. Alam, "Distributed intrusion detection system in a multi-layer network architecture of smart grids," *IEEE Transactions on Smart Grid*, vol. 2(4), pp. 796-808, 2011.
- [12] S. A. Zonouz, H. Khurana, W. H. Sanders, and T. M. Yardley, "RRE: A game-theoretic intrusion response and recovery engine," in *Proc. IEEE/IFIP International Conference on Dependable Systems & Networks (DSN'09)*, 2009, pp. 439-448.
- [13] C. A. Carver, *Intrusion Response Systems: A Survey*, Department of Computer Science, Texas A&M University, College Station, TX, 2000.